

Amendments to the Claims

1. (Currently Amended) A method of dynamically controlling the
2 rate of communication between two entities, the method comprising:
receiving an electronic communication, for a first channel between a first entity
4 and a second entity, at a relay element situated between the first entity and the second
entity;
6 retrieving from said communication a modifiable first value associated with a first
target bandwidth for said first channel;
8 retrieving from said communication a fixed second value associated with a
desired bandwidth for said first channel, wherein the desired bandwidth is never less than
10 said first target bandwidth;
determining whether said relay element can provide said first target bandwidth for
12 said first channel; and
modifying said first value in said communication to a value associated with a
14 decreased first target bandwidth if said relay element cannot provide said first target
bandwidth for said channel.

2. (Original) The method of claim 1, further comprising:
2 forwarding said communication;
wherein said first value in said forwarded communication indicates a bandwidth
4 allocated to said first channel by said relay element.

3. (Original) The method of claim 1, further comprising, prior to said
2 determining:
receiving a set of communications on a set of channels through said switching
4 element, not including said first channel;
retrieving from said set of communications a set of values associated with target
6 bandwidths for said set of channels; and
summing said target bandwidths to calculate a total allocated bandwidth for said
8 relay element.

4. (Original) The method of claim 3, wherein said determining
2 comprises:
comparing said total allocated bandwidth to a maximum bandwidth of said relay
4 element; and
if said maximum bandwidth exceeds said total allocated bandwidth by a
6 difference of more than said first target bandwidth, determining that said relay element
can provide said first target bandwidth for said first channel.

5. (Original) The method of claim 1, wherein said determining
2 comprises:
comparing said first target bandwidth for said first channel to a previous
4 bandwidth granted to said first channel by said relay element; and
if said first target bandwidth is greater than said previous bandwidth, comparing a
6 difference between said first target bandwidth and said previous bandwidth with an
unallocated bandwidth of said relay element.

6. (Original) The method of claim 1, wherein said modifying comprises
2 changing said first value to a value associated with zero bandwidth.

7. Cancelled

8. (Original) The method of claim 1, wherein said first value is a time
2 value representing a time between communication transmissions from the first entity to
the second entity on said first channel.

9. (Original) The method of claim 1, wherein said electronic
2 communication is a packet.

10. (Original) The method of claim 9, wherein said relay element is a
2 switch and wherein said first entity and said second entity are computer systems.

11. (Original) The method of claim 1, wherein one of said first entity and
2 said second entity is a computer system; and
wherein the other of said first entity and said second entity is an input/output
4 subsystem.

12. (Currently Amended) A computer readable storage medium
2 storing instructions that, when executed by a computer, cause the computer to perform a
method of dynamically controlling the rate of communication between two entities, the
4 method comprising:
receiving an electronic communication, for a first channel between a first entity
6 and a second entity, at a relay element situated between the first entity and the second
entity;
8 retrieving from said communication a modifiable first value associated with a first
target bandwidth for said first channel;
10 retrieving from said communication a fixed second value associated with a
desired bandwidth for said first channel, wherein the desired bandwidth is never less than
12 said first target bandwidth;
determining whether said relay element can provide said first target bandwidth for
14 said first channel; and
modifying said first value in said communication to a value associated with a
16 decreased first target bandwidth if said relay element cannot provide said first target
bandwidth for said channel.

13. (Currently Amended) A method of dynamically controlling the
2 rate of communication between two entities, comprising:
generating at a first entity a first electronic communication for transmission to a
4 second entity over a first communication channel, wherein said first communication
includes a first value indicating a target rate of communication for said channel;
6 receiving said first communication at a switching element;
~~determining whether said switching element can provide said target rate of~~

8 ~~communication for said first channel;~~
 determining whether a maximum rate of communication of said switching
10 element has been allocated;
 if said maximum rate has not been allocated, identifying an available rate of
12 communication of said switching element by:
 (a) receiving a communication prior to said first communication at
14 said switching element, on a channel other than said first channel;
 (b) allocating a portion of the maximum rate of communication of said
16 switching element to said other channel;
 (c) repeating said steps (a) - (b) for a predetermined period of time;
18 (d) summing said rates of communication allocated to said other
 channels to determine a total allocated rate of communication; and
20 (e) determining the difference between the maximum rate of
 communication and said total allocated rate of communication;
22 if said switching element cannot provide said target rate of communication,
altering said first value to indicate a lower target rate of communication for said first
24 channel;
 receiving said first communication at said second entity; and
26 communicating said first value to said first entity.

14. (Original) The method of claim 13, further comprising determining
2 whether said switching element previously allocated a rate of communication to said first
channel.

15. (Original) The method of claim 13, further comprising after said
2 communicating:
 transmitting one or more communications from said first entity toward said
4 second entity at said lower target rate of communication.

16. (Original) The method of claim 13, wherein said generating comprises
2 storing said first value in said first communication prior to transmitting it over said first

channel.

17. (Original) The method of claim 16, wherein said generating further
2 comprises storing a second value in said first communication; and
wherein said second value indicates a requested rate of communication for said
4 channel.

18. (Original) The method of claim 17, wherein said first value is equal to
2 said second value.

19. (Currently Amended) The method of claim 17, wherein one or
2 more of said first value and said second value comprises a threshold value indicating a
maximum rate of communication for said channel desired by the first entity.

20. (Currently Amended) The method of claim 17, wherein one or
2 more of said first value and said second value comprise a threshold value indicating a
minimum rate of communication for said channel desired by the first entity.

21. (Original) The method of claim 20, further comprising at said
2 switching element:
detecting said threshold value indicating said minimum rate of communication;
4 and
tearing down said channel.

22. (Original) The method of claim 17, wherein one or more of said first
2 value and said second value comprises a time period representing a delay between
transmission of successive communications over said first channel from said first entity;
4 and
wherein said rate of communication indicated by said time period is substantially
6 equal to the inverse of said time period.

23. Cancelled

24. Cancelled

25. Cancelled

26. (Original) The method of claim 13, wherein said altering comprises
2 setting said first value to a threshold value indicating a minimum rate of communication.

27. (Original) The method of claim 26, further comprising at said first
2 entity after said communicating:
4 ceasing transmission of communications to said second entity over said first
4 channel.

28. (Original) The method of claim 13, wherein said first value is a time
2 period between successive electronic communication transmissions from said first entity
on said first channel.

29. (Original) The method of claim 28, wherein said target rate of
2 communication is substantially equal to the inverse of said first value.

30. (Original) The method of claim 13, wherein said first value is a
2 measure of bandwidth.

31. (Currently Amended) A computer readable storage medium
2 storing instructions that, when executed by a computer, cause the computer to perform a
method of dynamically controlling the rate of communication between two entities, the
4 method comprising:
6 generating at a first entity a first electronic communication for transmission to a
second entity over a first communication channel, wherein said first communication
includes a first value indicating a target rate of communication for said channel;

8 receiving said first communication at a switching element;
determining whether said switching element can provide said target rate of
10 communication for said first channel;
determining whether a maximum rate of communication of said switching
12 element has been allocated;
if said maximum rate has not been allocated, identifying an available rate of
14 communication of said switching element by:
(a) receiving a communication prior to said first communication at
16 said switching element, on a channel other than said first channel;
(b) allocating a portion of the maximum rate of communication of said
18 switching element to said other channel;
(c) repeating said steps (a) - (b) for a predetermined period of time;
20 (d) summing said rates of communication allocated to said other
channels to determine a total allocated rate of communication; and
22 (e) determining the difference between the maximum rate of
communication and said total allocated rate of communication;
24 if said switching element cannot provide said target rate of communication,
altering said first value to indicate a lower target rate of communication for said first
26 channel;
receiving said first communication at said second entity; and
28 communicating said first value to said first entity.

32. (Currently Amended) A method of controlling a network
2 communication rate, the method comprising:
receiving at a downstream ~~an~~ intermediate node ~~coupling a first network node and~~
4 ~~a second network node~~ a fixed rate value representing a desired rate of communication
for a channel between a the first network node and a the second network node, and a
6 modifiable value representing a target rate of communication allocated to the channel by
an upstream intermediate node; and
8 at the downstream intermediate node, allocating to the channel a rate of
communication higher than the target rate of communication if the downstream

10 intermediate node has sufficient available bandwidth; and
12 if the downstream intermediate node does not have sufficient available bandwidth
14 to if the intermediate node cannot conduct communications between the first network
16 node and the second network node at on the channel at a rate greater than or equal to said
target rate value, adjusting decreasing said modifiable rate value such that the
intermediate node can conduct communications between the first network node and the
second network node at on the channel at an adjusted rate represented by said adjusted
modifiable rate value.

33. (Currently Amended) The method of claim 32, wherein each said
2 rate value corresponds to is a time between communications transmitted from the first
network node toward the second network node.

34. (Currently Amended) The method of claim 33, wherein said
2 adjusting decreasing comprises increasing said time between communications.

35. (Currently Amended) The method of claim 32, wherein if said
2 modifiable rate value is adjusted to a first threshold decreased to a first value, the first
network node stops sending communications toward the second network node through
4 the channel intermediate node.

36. (Currently Amended) The method of claim 32, wherein if said
2 modifiable rate value received at the downstream intermediate node is adjusted to a
second threshold has a second value, the first network node sends communications
4 toward the second network node through the channel intermediate node at a maximum
rate.

37. (Currently Amended) The method of claim 32, further comprising:
2 notifying the first network node of said adjusted modifiable decreased rate value;
wherein the first network node then transmits communications toward the second
4 network node through the channel at said adjusted decreased rate value.

38. Cancelled

39. Cancelled

40. (Currently Amended) The method of claim 32, wherein the
2 downstream intermediate node is InfiniBand compliant.

41. (Currently Amended) The method of claim 32, wherein the
2 downstream intermediate node is a switch.

42. (Currently Amended) The method of claim 32, wherein the
2 downstream intermediate node is a router.

43. (Currently Amended) The method of claim 32, wherein the
2 downstream intermediate node is a hub.

44. (Currently Amended) The method of claim 32, wherein the
2 downstream intermediate node is a bridge.

45. (Currently Amended) The method of claim 32, wherein the
2 downstream intermediate node is a repeater.

46. (Currently Amended) The method of claim 32, wherein the
2 downstream intermediate node is a network adapter.

47. (Currently Amended) The method of claim 32, wherein the
2 downstream intermediate node is a computer.

48. (Currently Amended) The method of claim 32, wherein the
2 downstream intermediate node is a communication bus.

49. (Currently Amended) A computer readable storage medium
2 storing instructions that, when executed by a computer, cause the computer to perform a
method of controlling a network communication rate, the method comprising:
4 receiving at a downstream ~~an~~ intermediate node ~~coupling a first network node and~~
~~a second network node~~ a fixed rate value representing a desired rate of communication
6 for a channel between ~~a~~ the first network node and ~~a~~ the second network node, and a
modifiable value representing a target rate of communication allocated to the channel by
8 an upstream intermediate node; and
at the downstream intermediate node, allocating to the channel a rate of
10 communication higher than the target rate of communication if the downstream
intermediate node has sufficient available bandwidth; and
12 if the downstream intermediate node does not have sufficient available bandwidth
to if the intermediate node cannot conduct communications ~~between the first network~~
14 ~~node and the second network node at~~ on the channel at a rate greater than or equal to said
target rate value, adjusting decreasing said modifiable rate value such that the
16 intermediate node can conduct communications ~~between the first network node and the~~
~~second network node at~~ on the channel at an adjusted rate represented by said adjusted
18 modifiable rate value.

50. Cancelled

51. Cancelled

52. (Currently Amended) A computer readable storage medium
2 containing a data structure configured to indicate a rate of communication over a
communication channel, the data structure comprising:
4 a header portion comprising:
an identifier of an originator of said data structure;
6 an identifier of a destination of said data structure; ~~and~~
a first value corresponding to a target bandwidth ~~rate of communication~~

8 between said originator and said destination; and
a second value corresponding to a requested bandwidth between said
10 originator and said destination;
wherein said second value is fixed; and
12 wherein said first value is modifiable during transmission of said data structure
from said originator to said destination but the target bandwidth represented by said first
14 value can never be greater than said requested bandwidth.

53. (Original) The computer readable storage medium of claim 52,
2 wherein said first value of said header portion of said data structure comprises a time
period and said target rate of communication is substantially equal to the inverse of said
4 time period.

54. (Original) The computer readable storage medium of claim 52, said
2 data structure further comprising:
a data portion comprising a set of data.

55. Cancelled

56. Cancelled

57. (Currently Amended) An apparatus for dynamically adjusting the
2 rate of communications between a first entity and a second entity on a channel,
comprising:
4 a communication port configured to forward a communication received from a
first entity toward a second entity on a communication channel;
6 a first memory configured to store said communication;
a second memory configured to store a target bandwidth for said channel, wherein
8 said target bandwidth is indicated by a modifiable first value in said communication;
a third memory configured to store a requested bandwidth for said channel,
10 wherein said requested bandwidth is indicated by a fixed second value in said

communication and is never less than said target bandwidth;

12 a comparator configured to compare one or more of said target bandwidth and
said requested bandwidth to an available bandwidth for said port; and

14 a processor configured to:

allocate to said channel a bandwidth equal to or greater than said target
16 bandwidth, up to said requested bandwidth, if the available bandwidth is
 sufficient; and

18 adjust said first value to indicate a different target bandwidth if the
 available bandwidth is insufficient to allow a bandwidth equal to or greater than
20 said target bandwidth to be allocated to said channel;

wherein said target bandwidth indicated by said first value received in said
22 communication is the bandwidth allocated to said channel upstream of said port.

58. (Currently Amended) The apparatus of claim 57, further
2 comprising an extractor configured to extract said first value and said second value from
said communication.

59. (Currently Amended) The apparatus of claim 58, wherein each of
2 said first value and said second value comprises a time period representing a delay
between communication transmissions from said first entity toward said second entity on
4 said channel, the apparatus further comprising:
 an inverter configured to invert said time period;
6 ~~wherein said target bandwidth is substantially equal to said inverted time period.~~

60. (Currently Amended) The apparatus of claim 59, further
2 comprising:
 an adder configured to add said target bandwidth indicated by said first value of
4 said communication to a target bandwidth indicated by a value within of a previous
communication on a different channel to calculate a total ~~allocated~~ bandwidth allocated
6 by said port.

61. (Original) The apparatus of claim 60, wherein said available
2 bandwidth is substantially equal to a maximum bandwidth of said port minus said total
allocated bandwidth.

62. Cancelled

63. (Currently Amended) The apparatus of claim 57, wherein said
2 processor is configured to adjust said first value to indicate a lower target bandwidth if
said apparatus is unable to provide said target bandwidth or a higher bandwidth.

64. Cancelled

65. (New) The method of claim 1, wherein said modifying comprises
2 replacing said modifiable first value with a modified first value associated with a lower
target bandwidth, the method further comprising:
4 at said relay element, allocating the lower target bandwidth to said first channel;
and
6 at another relay element downstream of said relay element:
receiving said electronic communication containing said modified first
8 value and said fixed second value; and
allocating to said first channel a bandwidth higher than the lower target
10 bandwidth.